

CLAIMS

1. A method for managing a data storage system that includes first and second storage media, the method comprising:

5 maintaining a record predictive of locations to which data are to be written on the first storage medium by a host computer;

receiving a write command from the host computer directed to storing specified data at a specified
10 location on the first storage medium;

if the specified location is not included in the record, updating the record responsively to the specified location;

signaling the host computer that the specified data
15 have been stored in the data storage system responsively to storing the specified data and, if the specified location was not included in the record, responsively to updating the record;

copying the specified data to the second storage
20 medium responsively to the record; and

storing the specified data in the specified location on both the first and second storage media.

2. The method according to claim 1, wherein copying the
25 specified data comprises transmitting the specified data between mutually-remote sites over a communication link between the sites.

3. The method according to claim 1, wherein copying the
30 specified data comprises asynchronously mirroring on the

second storage medium the data that are stored on the first storage medium.

4. The method according to claim 1, wherein maintaining
5 and updating the record comprise marking respective bits in a bitmap corresponding to the locations to which the data are to be written on the first and second non-volatile storage media.

10 5. The method according to claim 1, wherein updating the record comprises predicting one or more further locations to which the host computer is expected to write the data in a subsequent write operation, and adding the one or more further locations to the record.

15 6. The method according to claim 1, wherein updating the record comprises removing one or more locations, other than the specified location, from the record, so as to limit a size of the record.

20 7. The method according to claim 1, wherein maintaining the record comprises storing the record in a non-volatile memory, and wherein updating the record comprises modifying the record that is stored in the non-volatile
25 memory.

8. The method according to claim 7, wherein modifying the record comprises:
comparing the specified location to a copy of the
30 record held in a volatile memory;

modifying the copy of the record so that at least the specified location is included in the copy of the record; and

destaging the modified copy of the record to the
5 non-volatile memory.

9. The method according to claim 8, wherein the record is not modified in the non-volatile memory responsively to receiving the write command as long as the specified
10 location to which the specified data are to be written is included in the record.

10. The method according to claim 7, wherein modifying the record comprises adding a plurality of locations,
15 including the specified location, to the record.

11. The method according to claim 7, wherein receiving the write command comprises intercepting the write command using a data storage appliance, which comprises
20 the non-volatile memory, wherein the data storage appliance copies the specified data to the second storage medium, and maintains and updates the record in the non-volatile memory.

25 12. The method according to claim 11, wherein copying the specified data comprises, upon recovery from a failure of the data storage appliance, reading the record from the non-volatile memory, and copying the data to the second storage medium from the locations on the first
30 storage medium that are indicated in the record.

13. The method according to claim 1, wherein receiving the write command comprises intercepting the write command using a data storage appliance, which maintains and updates the record, and which copies the specified data to the second storage medium.

14. The method according to claim 13, wherein copying the specified data comprises copying the data, using the data storage appliance, from the specified location on the first storage medium to the second storage medium.

15. The method according to claim 14, wherein copying the data comprises selecting the locations on the first storage medium, including the specified location, from which to copy the data responsively to the record.

16. The method according to claim 13, wherein intercepting the write command comprises receiving the write command at a first data storage appliance, and wherein maintaining the record comprises holding the record on a second data storage appliance, and wherein updating the record comprises sending a message from the first data storage appliance to the second data storage appliance so as to cause the second data storage appliance to update the record.

17. The method according to claim 16, wherein updating the record further comprises receiving an acknowledgment of the message at the first data storage appliance from

the second data storage appliance, and wherein signaling the host computer comprises notifying the host computer that the specified data have been stored responsively to receiving the acknowledgment.

5

18. The method according to claim 16, and comprising, upon occurrence of a failure in the first data storage appliance, copying the data to the second storage medium using the second data storage appliance from the
10 locations on the first storage medium that are indicated in the record.

19. The method according to claim 16, wherein maintaining the record comprises maintaining on the first
15 data storage appliance a copy of the record held on the second data storage appliance, and wherein sending the message comprises deciding at the first data storage appliance to send the message responsively to the copy of the record.

20

20. The method according to claim 19, wherein sending the message comprises modifying both the record and the copy of the record responsively to the specified location.

25

21. The method according to claim 20, wherein modifying both the record and the copy of the record comprises adding a plurality of locations, including the specified location, to both the record and the copy of the record.

30

22. The method according to claim 21, wherein
maintaining the copy of the record comprises selecting
one or more locations, other than the specified location,
to be removed from the record, and instructing the second
5 data storage appliance to remove the one or more
locations from the record, so as to limit a size of the
record.

23. A method for managing a data storage system that
10 includes a local storage medium, the method comprising:

coupling first and second data storage appliances to
provide one or more host computers with access to the
local storage medium via the data storage appliances;

maintaining on the second data storage appliance a
15 record indicative of first locations on the local storage
medium to which a host computer has directed a write
command via the first data storage appliance;

adding to the record one or more entries that are
predictive of second locations on the local storage
20 medium to which data are expected to be written via the
first data storage appliance;

receiving at the second data storage appliance, from
one of the host computers, a request to access the data
stored at a specified location on the local storage
25 medium; and

if the specified location is included in the record,
sending a message from the second data storage appliance
to the first data storage appliance so as to ascertain a
status of the data at the specified location before
30 permitting the one of the host computers to access the
data stored at the specified location.

24. The method according to claim 23, wherein coupling the first and second data storage appliances comprises configuring the data storage appliances to virtualize the access by the host computers to the local storage medium.

5

25. The method according to claim 23, and comprising copying the data stored on the local storage medium to a remote storage medium, using at least the first data storage appliance, in a data mirroring process.

10

26. The method according to claim 25, wherein copying the data comprises using both the first and second data storage appliances in the data mirroring process, wherein upon occurrence of a failure in the first data storage
15 appliance, the second data storage appliance continues copying the data from the first and second locations responsively to the record.

27. The method according to claim 25, wherein the data
20 mirroring process comprises an asynchronous data mirroring processing, and wherein maintaining the record comprises removing the locations from the record after the data stored at the locations have been copied to the a remote storage medium.

25

28. The method according to claim 23, and comprising permitting the one of the host computers to access the data stored at the specified location via the second data storage appliance in response to the request without

sending the message if the specified location is not included in the record.

29. The method according to claim 23, wherein adding the
5 one or more entries comprises, upon receiving the write
command at the first data storage appliance to write the
data to a given location, sending an instruction from the
first data storage appliance to the second data storage
10 appliance, which causes the second data storage appliance
to include the given location and one or more of the
second locations in the record.

30. The method according to claim 29, wherein
maintaining the record comprises maintaining a copy of
15 the record on the first data storage appliance, and
wherein sending the instruction comprises determining
whether to send the instruction responsively to the copy
of the record.

20 31. The method according to claim 30, wherein
determining whether to send the instruction comprises
deciding to send the instruction only if the given
location is not included in the record.

25 32. The method according to claim 23, wherein receiving
the request from the one of the host computers comprises
receiving a further write command from the one of the
host computers at the second data storage appliance, and
updating the data stored at the specified location, using

the second data storage appliance, responsively to the further write command.

33. The method according to claim 32, and comprising:

5 maintaining a further record on the first data storage appliance, indicative of third locations at the which the data have been updated by the second data storage appliance and including further entries that are predictive of fourth locations on the local storage
10 medium to which data are expected to be written via the second data storage appliance;

 receiving at the first data storage appliance, from a further one of the host computers, a further request to access the data stored at a further location on the local
15 storage medium; and

 if the further location is included in the record, sending a further message from the first data storage appliance to the second data storage appliance so as to ascertain the status of the data at the further location
20 before permitting the further one of the host computers to access the data stored at the further location.

34. The method according to claim 23, and comprising, upon occurrence of a failure in the second data storage
25 appliance, configuring the first data storage appliance to provide the data directly to the one of the host computers.

35. The method according to claim 23, wherein
30 maintaining the record comprises marking respective bits

in a bitmap corresponding to the locations that are included in the record.

36. Apparatus for data storage, comprising:

5 first and second storage media; and
at least one data storage appliance, comprising a memory and a control unit, which is coupled to intercept write commands issued by a host computer for writing data to the first storage medium, and to copy the data to the
10 second storage medium in a mirroring process so that the data are stored on both the first and second storage media, and

wherein the control unit is arranged to maintain in the memory a record predictive of locations to which the
15 data are to be written on the first storage medium by a host computer, and to copy the data to the second storage medium responsively to the record.

37. The apparatus according to claim 36, wherein the
20 first and second storage media are located at mutually-remote sites, and wherein the control unit is arranged to transmit the data to the second storage media over a communication link between the sites.

25 38. The apparatus according to claim 36, wherein the mirroring process comprises an asynchronous mirroring process.

39. The apparatus according to claim 36, wherein the
30 record in the memory comprises a bitmap.

40. The apparatus according to claim 36, wherein the control unit is further arranged, upon receiving a write command from the host computer directed to storing
5 specified data at a specified location on the first storage medium, to update the record responsively to the specified location if the specified location is not included in the record, and to signal the host computer that the specified data have been stored in the data
10 storage system responsively to storing the specified data and, if the specified location was not included in the record, responsively to updating the record

41. The apparatus according to claim 40, wherein the
15 control unit is arranged to update the record by predicting one or more further locations to which the host computer is expected to write the data in a subsequent write operation, and adding the one or more further locations to the record.

20

42. The apparatus according to claim 40, wherein the control unit is further arranged to update the record by removing one or more locations, other than the specified location, from the record, so as to limit a size of the
25 record.

43. The apparatus according to claim 40, wherein the memory of the at least one data storage appliance comprises a non-volatile memory, and wherein the control
30 unit is arranged to store the record in the non-volatile

memory, and to update the record by modifying the record that is stored in the non-volatile memory.

44. The apparatus according to claim 43, wherein the
5 memory of the at least one data storage appliance further
comprises a volatile memory, and wherein the control unit
is arranged to hold a copy of the record in the volatile
memory, and to modify the record by comparing the
10 specified location to the copy of the record, modifying
the copy of the record so that at least the specified
location is included in the copy of the record, and
destaging the modified copy of the record to the non-
volatile memory.

15 45. The apparatus according to claim 44, wherein the
record is not modified in the non-volatile memory
responsively to receiving the write command as long as
the specified location to which the specified data are to
be written is included in the record.

20

46. The apparatus according to claim 43, wherein the
control unit is arranged, upon modifying the record, to
add a plurality of locations, including the specified
location, to the record.

25

47. The apparatus according to claim 43, wherein the
control unit is arranged, upon recovery from a failure of
the at least one data storage appliance, to read the
record from the non-volatile memory, and to copy the data

to the second storage medium from the locations on the first storage medium that are indicated in the record.

48. The apparatus according to claim 40, wherein the at
5 least one data storage appliance comprises first and
second data storage appliances, wherein the record is
held on the second data storage appliance, and wherein
the first data storage appliance is arranged, upon
receiving the write command from the host computer, to
10 send a message to the second data storage appliance so as
to cause the second data storage appliance to update the
record.

49. The apparatus according to claim 48, wherein the
15 second data storage appliance is arranged to send an
acknowledgment of the message to the first data storage
appliance, and wherein the first data storage appliance
is arranged to notify the host computer that the
specified data have been stored responsively to receiving
20 the acknowledgment.

50. The apparatus according to claim 48, wherein the
second data storage appliance is arranged, upon
occurrence of a failure in the first data storage
25 appliance, to copy the data to the second storage medium
from the locations on the first storage medium that are
indicated in the record.

51. The apparatus according to claim 48, wherein the
30 first data storage appliance is arranged to maintain a

copy of the record held on the second data storage appliance, and to decide to send the message responsively to the copy of the record.

5 52. The apparatus according to claim 51, wherein both the record and the copy of the record are modified responsively to the message.

10 53. The apparatus according to claim 52, wherein the first and second data storage appliances are arranged to modify the record and the copy of the record by adding a plurality of locations, including the specified location, to both the record and the copy of the record.

15 54. The apparatus according to claim 48, wherein the first data storage appliance is arranged to select one or more locations, other than the specified location, to be removed from the record, and to instruct the second data storage appliance to remove the one or more locations
20 from the record, so as to limit a size of the record.

55. Data storage apparatus, comprising:
a local storage medium; and
first and second data storage appliances, which are
25 coupled to provide one or more host computers with access to the local storage medium via the data storage appliances,

wherein the second data storage appliance is arranged to maintain a record indicative of first
30 locations on the local storage medium to which a host

computer has directed a write command via the first data storage appliance, and to add to the record one or more entries that are predictive of second locations on the local storage medium to which data are expected to be
5 written via the first data storage appliance, and

wherein the second data storage appliance is arranged, upon receiving a request from one of the host computers to access the data stored at a specified location on the local storage medium, if the specified
10 location is included in the record, to send a message to the first data storage appliance so as to ascertain a status of the data at the specified location before permitting the one of the host computers to access the data stored at the specified location.

15 56. The apparatus according to claim 55, wherein the first and second data storage appliances are configured to virtualize the access by the host computers to the local storage medium.

20 57. The apparatus according to claim 55, wherein at least the first data storage appliance is arranged to copy the data stored on the local storage medium to a remote storage medium in a data mirroring process.

25 58. The apparatus according to claim 57, wherein the second data storage appliance is arranged, upon occurrence of a failure in the first data storage appliance, to continue copying the data from the first and second locations to the remote storage medium responsively to the record.

30 59. The apparatus according to claim 57, wherein the data mirroring process comprises an asynchronous data mirroring processing, and wherein the data storage

appliances are arranged to remove the locations from the record after the data stored at the locations have been copied to the a remote storage medium.

60. The apparatus according to claim 55, wherein the
5 second data storage appliance is arranged to permit the one of the host computers to access the data stored at the specified location via the second data storage appliance in response to the request without sending the message if the specified location is not included in the
10 record.

61. The apparatus according to claim 55, wherein the first data storage appliance is arranged, upon receiving the write command from the host computer to write the data to a given location, to send an instruction to the
15 second data storage appliance, which causes the second data storage appliance to include the given location and one or more of the second locations in the record.

62. The apparatus according to claim 61, wherein the first data storage appliance is arranged to maintain a
20 copy of the record, and to determine whether to send the instruction responsively to the copy of the record.

63. The apparatus according to claim 62, wherein the first data storage appliance is arranged to send the instruction only if the given location is not included in
25 the record.

64. The apparatus according to claim 55, wherein the second data storage appliance is arranged to receive a further write command from the one of the host computers, and to update the data stored at the specified location
30 responsively to the further write command and to the record.

65. The apparatus according to claim 64, wherein the first data storage appliance is arranged to maintain a further record, indicative of third locations at the
5 which the data have been updated by the second data storage appliance and including further entries that are predictive of fourth locations on the local storage medium to which data are expected to be written via the second data storage appliance, and

10 wherein the first data storage appliance is arranged, upon receiving from a further one of the host computers a further request to access the data stored at a further location on the local storage medium, if the further location is included in the record, to send a
15 further message to the second data storage appliance so as to ascertain the status of the data at the further location before permitting the further one of the host computers to access the data stored at the further location.

20

66. The apparatus according to claim 55, wherein the first data storage appliance is arranged, upon occurrence of a failure in the second data storage appliance, to provide the data directly to the one of the host
25 computers.

67. The apparatus according to claim 55, wherein the record comprises a bitmap.

68. A computer software product for use in a data storage system that includes first and second storage media and at least one data storage appliance having a control unit, the product comprising a computer-readable medium in which program instructions are stored, which instructions, when read by the control unit, cause the data storage appliance to intercept write commands issued by a host computer for writing data to the first storage medium, and to copy the data to the second storage medium in a mirroring process so that the data are stored on both the first and second storage media,

wherein the instructions cause the control unit to maintain a record predictive of locations to which the data are to be written on the first storage medium by a host computer, and to copy the data to the second storage medium responsively to the record.

69. The product according to claim 68, wherein the first and second storage media are located at mutually-remote sites, and wherein the instructions cause the control unit to transmit the data to the second storage media over a communication link between the sites.

70. The product according to claim 68, wherein the mirroring process comprises an asynchronous mirroring process.

71. The product according to claim 68, wherein the record comprises a bitmap.

30

72. The product according to claim 68, wherein the instructions cause the control unit, upon receiving a write command from the host computer directed to storing specified data at a specified location on the first storage medium, to update the record responsively to the specified location if the specified location is not included in the record, and to signal the host computer that the specified data have been stored in the data storage system responsively to storing the specified data and, if the specified location was not included in the record, responsively to updating the record

73. The product according to claim 72, wherein the instructions cause the control unit to update the record by predicting one or more further locations to which the host computer is expected to write the data in a subsequent write operation, and adding the one or more further locations to the record.

74. The product according to claim 72, wherein the instructions cause the control unit to update the record by removing one or more locations, other than the specified location, from the record, so as to limit a size of the record.

75. The product according to claim 72, wherein the at least one data storage appliance comprises a non-volatile memory, and wherein the instructions cause the control unit to store the record in the non-volatile memory, and to update the record by modifying the record that is stored in the non-volatile memory.

76. The product according to claim 75, wherein the at least one data storage appliance further comprises a volatile memory, and wherein the instructions cause the control unit to hold a copy of the record in the volatile memory, and to modify the record by comparing the specified location to the copy of the record, modifying the copy of the record so that at least the specified location is included in the copy of the record, and destaging the modified copy of the record to the non-volatile memory.

77. The product according to claim 76, wherein the record is not modified in the non-volatile memory responsively to receiving the write command as long as the specified location to which the specified data are to be written is included in the record.

78. The product according to claim 75, wherein the instructions cause the control unit, upon modifying the record, to add a plurality of locations, including the specified location, to the record.

79. The product according to claim 75, wherein the instructions cause the control unit, upon recovery from a failure of the at least one data storage appliance, to read the record from the non-volatile memory, and to copy the data to the second storage medium from the locations on the first storage medium that are indicated in the record.

80. The product according to claim 72, wherein the at least one data storage appliance comprises first and second data storage appliances, which comprise respective
5 first and second control units, wherein the record is held on the second data storage appliance, and wherein the instructions cause the first control unit, upon receiving the write command from the host computer, to send a message to the second data storage appliance so as
10 to cause the second data storage appliance to update the record.

81. The product according to claim 80, wherein the instructions cause the second control unit to send an
15 acknowledgment of the message to the first data storage appliance, and cause the first control unit to notify the host computer that the specified data have been stored responsively to receiving the acknowledgment.

20 82. The product according to claim 80, wherein the instructions cause the second control unit, upon occurrence of a failure in the first data storage appliance, to copy the data to the second storage medium from the locations on the first storage medium that are
25 indicated in the record.

83. The product according to claim 80, wherein the instructions cause the first control unit to maintain a copy of the record held on the second data storage

appliance, and to decide to send the message responsively to the copy of the record.

84. The product according to claim 83, wherein both the
5 record and the copy of the record are modified
responsively to the message.

85. The product according to claim 84, wherein the
instructions cause the first and second control units to
10 modify the record and the copy of the record by adding a
plurality of locations, including the specified location,
to both the record and the copy of the record.

86. The product according to claim 80, wherein the
15 instructions cause the first control unit to select one
or more locations, other than the specified location, to
be removed from the record, and to instruct the second
data storage appliance to remove the one or more
locations from the record, so as to limit a size of the
20 record.

87. A computer software product for use in a data
storage system that includes a local storage media and
first and second data storage appliances, which include
25 respective first and second control units and are coupled
to provide one or more host computers with access to the
local storage medium via the data storage appliances, the
product comprising a computer-readable medium in which
program instructions are stored, which instructions, when
30 read by the control units, cause the second control unit

to maintain a record indicative of first locations on the local storage medium to which a host computer has directed a write command via the first data storage appliance, and to add to the record one or more entries
5 that are predictive of second locations on the local storage medium to which data are expected to be written via the first data storage appliance, and

wherein the instructions further cause the second control unit, upon receiving a request from one of the
10 host computers to access the data stored at a specified location on the local storage medium, if the specified location is included in the record, to send a message to the first data storage appliance so as to ascertain a status of the data at the specified location before
15 permitting the one of the host computers to access the data stored at the specified location.

88. The product according to claim 87, wherein the instructions cause the first and second control units to
20 virtualize the access by the host computers to the local storage medium.

89. The product according to claim 87, wherein the instructions cause at least the first control unit to
25 copy the data stored on the local storage medium to a remote storage medium in a data mirroring process.

90. The product according to claim 89, wherein the instructions cause the second control unit, upon
30 occurrence of a failure in the first data storage appliance, to continue copying the data from the first

and second locations to the remote storage medium
responsively to the record.

91. The product according to claim 89, wherein the data
5 mirroring process comprises an asynchronous data
mirroring processing, and wherein the instructions cause
the second control unit to remove the locations from the
record after the data stored at the locations have been
copied to the a remote storage medium.

10

92. The product according to claim 87, wherein the
instructions cause the second control unit to permit the
one of the host computers to access the data stored at
the specified location via the second data storage
15 appliance in response to the request without sending the
message if the specified location is not included in the
record.

93. The product according to claim 87, wherein the
20 instructions cause the first control unit, upon receiving
the write command from the host computer to write the
data to a given location, to send an instruction to the
second data storage appliance, which causes the second
data storage appliance to include the given location and
25 one or more of the second locations in the record.

94. The product according to claim 93, wherein the
instructions cause the first control unit to maintain a
copy of the record, and to determine whether to send the
30 instruction responsively to the copy of the record.

95. The product according to claim 94, wherein the instructions cause the first control unit to send the instruction only if the given location is not included in
5 the record.

96. The product according to claim 87, wherein the instructions cause the second control unit, upon receiving a further write command from the one of the
10 host computers, to update the data stored at the specified location responsively to the further write command and to the record.

97. The product according to claim 96, wherein the
15 instructions cause the first control unit to maintain a further record, indicative of third locations at the which the data have been updated by the second data storage appliance and including further entries that are predictive of fourth locations on the local storage
20 medium to which data are expected to be written via the second data storage appliance, and

wherein the instructions further cause the first control unit, upon receiving from a further one of the host computers a further request to access the data
25 stored at a further location on the local storage medium, if the further location is included in the record, to send a further message to the second data storage appliance so as to ascertain the status of the data at the further location before permitting the further one of
30 the host computers to access the data stored at the further location.

98. The product according to claim 87, wherein the instructions cause the first control unit, upon occurrence of a failure in the second data storage
5 appliance, to provide the data directly to the one of the host computers.

99. The product according to claim 87, wherein the record comprises a bitmap.